arXiv:1609.02121v2 [cs.SI] 23 Mar 2017

Generating realistic scaled complex networks *

Christian L. Staudt[2], Michael Hamann[2], Alexander Gutfraind[1], Ilya Safro[3], and Henning Meyerhenke[2]

[1]Laboratory for Mathematical Analysis of Complexity and Conflicts, Chicago, University of Illinois at Chicago, Chicago, IL, agutfraind.research@gmail.com
[2]Institute of Theoretical Informatics, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany,
{michael.hamann,meyerhenke,christian.staudt}@kit.edu
[3]School of Computing, Clemson University, Clemson, SC, isafro@clemson.edu

Abstract

Research on generative models is a central project in the emerging field of network science, and it studies how statistical patterns found in real networks could be generated by formal rules. Output from these generative models is then the basis for designing and evaluating computational methods on networks, and for verification and simulation studies. During the last two decades, a variety of models has been proposed with an ultimate goal of achieving comprehensive realism for the generated networks. In this study, we (a) introduce a new generator, termed ReCoN; (b) explore how ReCoN and some existing models can be fitted to an original network to produce a structurally similar replica, (c) use ReCoN to produce networks much larger than the original exemplar, and finally (d) discuss open problems and promising research directions. In a comparative experimental study, we find that ReCoN is often superior to many other state-of-the-art network generation methods. We argue that ReCoN is a scalable and effective tool for modeling a given network while preserving important properties at both micro- and macroscopic scales, and for scaling the exemplar data by orders of magnitude in size.

Keywords network generation · multiscale modeling · network modeling · communities

1 Introduction

Networks are widely used to represent connections between entities, because they provide intuitive windows into the function, dynamics, and evolution of natural and man-made systems. However, high-quality, large-scale network data is often unavailable because of economic, legal, technological, or other obstacles [17, 12]. For example, the human contact networks along with infectious diseases spread are notoriously difficult to estimate, and thus our understanding of the dynamics and control of epidemics stems from models that make highly simplifying assumptions or simulate contact networks from incomplete or proxy data [23, 32, 49]. In another domain, the development of cybersecurity systems requires testing across diverse threat scenarios and validation across diverse network structures that are not yet known, in anticipation of the computer networks of the future [21]. In both examples, the systems of interest cannot be represented by a single exemplar network, but must instead be modeled as collections of networks in which the variation among them may be just as important as their common features. Such cases point to the importance of data-driven methods for synthesizing networks that capture both the essential features of a system and realistic variability in order to use them in such tasks as simulations, analysis, and decision making.

*This work is partially supported by the DFG under grants ME 3619/3-1 and WA 654/22-1 within the Priority Program 1736 *Algorithms for Big Data*, and by the NSF awards #1522751 and #1647361. Funding was also provided by *Karlsruhe House of Young Scientists* via the *International Collaboration Package*. A preliminary version of the paper was presented at the 5th International Workshop on Complex Networks and their Applications [64].

A good network generator must meet two primary criteria: realism and diversity. The first, realism, needs to consider any properties of the network that govern the domain-specific processes of interest such as system function, dynamics, and evolution. Hence, realism may depend on both structural network features and the more subtle emerging features of the network. For example:

- Models of social networks should be able not only to reproduce structural features such as small-world properties, but also, and perhaps more importantly, to emulate emergent sociological phenomena such as interactions between individuals in a community, as driven by their psychological needs and daily routines. That is, the generated network should show similar interactions by its artificial individuals, as determined by implicit psychological and social rules.
- Models of connected solar energy collectors of different sizes and capacities should simulate realistic energy outputs influenced by the weather.
- Models of metabolic interactions should ultimately reflect biochemical properties of a cell.

Second, a synthetic network should reflect naturally occurring stochasticity in a system, without systematic bias that departs from reality. This feature is important for benchmarking and evaluating the robustness of network-based algorithms, anonymizing networks, and generating plausible hypothetical scenarios. In particular, when engineering algorithms, the ability to create good synthetic test data sets is valuable to estimate effectiveness and scalability of the proposed methods.

In addition, a network generator should be effective in tasks such as obfuscation (replacing restricted real data with similar synthetic data), compression (storing only a generator and its parameters instead of large graphs), as well as extrapolation and sampling (generating data at larger or smaller scales). Finally, the running time and memory requirements of the generator should be acceptable for realistically large datasets - datasets that may include millions or even billions of nodes and edges.

1.1 Problem definition

We envision the following usage scenarios. Given an original (or real) network $O = (V, E)$ ($n_o = |V|$, and $m_o = |E|$) that cannot be freely shared, we would like to be able to create a synthetic network R (with n_r nodes) that matches the original in essential structural properties, so that computational results obtained from processing this network are representative for what the original network would yield. We refer to R as a *replica*. We assume that whoever creates the replica has access to O and can pass it to a *model fitting* algorithm which uses it to parametrize a generative model.

More importantly, in addition to producing *scale-1 replicas* (where $n_r = n_o$), in the second scenario we want to use the generative model for *extrapolation*, i.e., we would like to parametrize it to produce produce a *scaled replica* R^x that has $n_r = x \cdot n_o$ nodes, where x is called the *scaling factor*. The structural properties of R^x should be such that they resemble a later growth stage of the original (also see Section 2). This should enable users of the replica to extrapolate the behavior of their methods when the network data is significantly scaled.

Finally, with respect to performance, we would like the generator algorithm and implementation as well as the fitting scheme to be efficient enough to produce large data sets (on the order of several millions of nodes and edges) quickly in practice.

1.2 Related work

A number of network generation methods have been developed, and these fall into two classes: generative models and editing methods (see surveys in [27, 17, 12, 21]). The first set of methods produces (by using such elementary operations as randomization and replication) networks from small initial seed networks (sometimes empty). The goal of such generation is to produce the structure that matches real data in *prespecified* properties, such as the degree distribution [2, 53, 45], clustering [5], and the number of small subgraphs. These methods are attractive because they often produce networks with the desired features and are grounded in well-developed theory (e.g., [22]). Some of these graph generators mechanistically model network growth [6, 34, 40], whereas

others incorporate domain-specific information such as geographic location [72] and cyber networks topological properties [46, 21].

One of the most successful generative strategies is based on Kronecker graphs [39] (including stochastic Kronecker graphs; see also related work [17, 55]). Graphs generated by this model preserve properties such as degree distribution, diameter, and some eigenvalues. Generative methods often describe an evolutionary process that can potentially lead to the original network; however, the probability that it will lead to the structure that is approximately isomorphic to the original one at some coarse-grained resolutions (but not necessarily at finest) is usually negligible. This makes generative methods ill-suited for studies such as simulations when one may need to work with systems that are similar to the original. Several generative models admit fast generators and are thus in our focus. Among those models are RMAT [18], BTER [33], and Hyperbolic Unit Disk Graphs (HUDG) [35]. However, a previous fitting scheme by Leskovec et al. [41] for RMAT graphs is quite time-consuming already for medium-sized networks [63].

The other class of network generators, graph editing [50], starts with a given (real or empirical) network and randomly changes its components until the network becomes sufficiently different than the original network. These are designed to introduce variability while preserving key structural properties. A multiscale network generator MUSKETEER [29] is an editing approach that is able to produce highly realistic synthetic networks by applying a local editing at different coarse-grained resolutions. The approach is comparable to ReCoN in the quality of generated networks; however, its implementation [28] requires further work to be applicable on very large networks.

1.3 Outline and contribution

In this paper we develop and evaluate a fast generator that focuses on creating realistic *scaled* replicas of complex networks. We point out in Section 2 which criteria we consider important for calling a (scaled) replica realistic. In particular we conceptualize realism in two ways: (i) matching an original graph in a set of important structural properties, and (ii) matching the running time behavior of various graph algorithms. Typically, being an important applied criteria, the last is not considered by most existing network generators.

Our new generator ReCoN, short for *Replication of Complex Networks* and described in Section 3, uses and extends ideas of LFR, a generator used for benchmarking community detection algorithms. Using the original degrees and a discovered community structure, we are able to capture a much more detailed signature of the network than a parametrization of the LFR generator. In Section 3.1 we discuss the generative models that we use for comparison (RMAT, HUDG, and BTER are among them) and develop model fitting schemes for them.

Our comparative experimental study in Section 4 indicates that ReCoN performs overall quite well and usually better than other generators in terms of realism. We can also conclude that the ReCoN implementation is fast, as it is capable of creating realistic scaled replicas on the scale of 10^8 edges in minutes. The ReCoN code is publicly available in the open-source network analysis package NetworKit [66]. In Section 5, we discuss open problems and promising research directions in generating realistic replicas of complex networks. This paper is based on the conference paper [64].

2 Realistic Replicas

We consider a generative model realistic if there is high structural similarity between the synthetic graphs produced and relevant real-world networks. It is neither our goal nor generally desirable to obtain an exact correspondence between original and replica. First, this would exclude the use case of obfuscation. Secondly, obtaining an isomorphic graph is rarely required for generalizable experiments. Note that we consider a single "realism score" for each model inappropriately reductionist. Rather, we quantify diverse aspects of realism in our experimental evaluation and leave it to the reader to decide about their relative importance.

For 1-scale replicas (with the same size as the original), we measure the similarity in terms of a set of commonly used metrics: Sparsity (number of edges vs number of nodes); degree distribution (more precisely its Gini coefficient); maximum degree as a proxy for the connectedness of hub

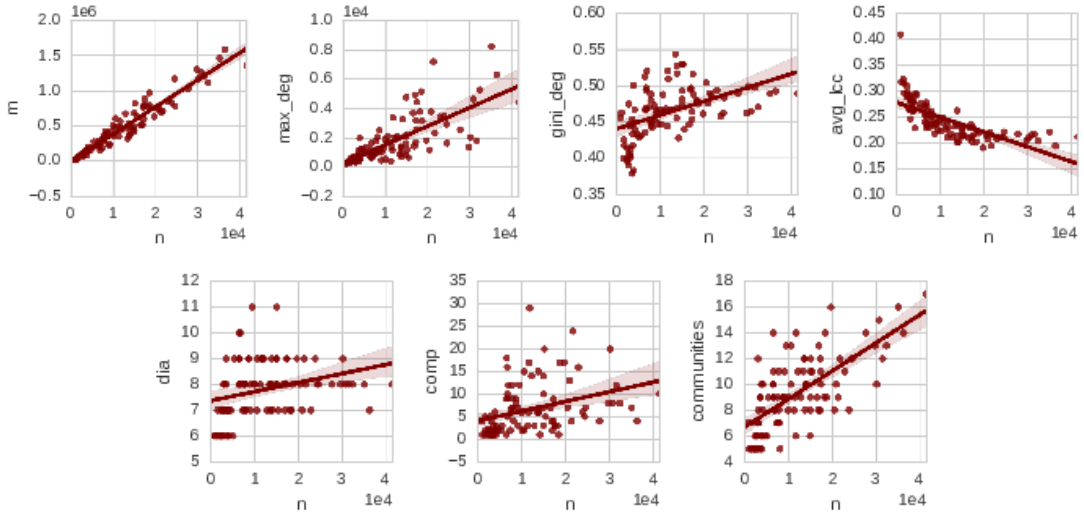

Figure 1: Scaling behavior of 100 Facebook networks; from left to right and top to bottom: number of edges, maximum degree, Gini coefficient of degree distribution, average local clustering coefficient, diameter, number of components, number of communities found by PLM

nodes; average local clustering coefficient to measure the local presence of triangles; diameter to monitor the small-world effect; number of connected components and number of communities as additional non-local features. These metrics cover both local and global properties and are deemed important characteristics of networks [54].

How can we extend the notion above regarding realism to *scaled* replicas of a network? To answer this question, let us look at the scaling behavior of a set of 100 Facebook social networks [68]. These networks were collected at an early stage of the Facebook online social networking service in which networks were still separated by universities. In Figure 1 we demonstrate basic structural measures of these Facebook networks against the number of nodes n, as well as a regression line and confidence intervals (shaded area) to emphasize the trend. While linear regression may not always seem completely appropriate for these data, the general trend is typically still captured.

We can observe from Figure 1 a growth of the number of edges m that is linear in n, an increase in the skew of the node degree distribution as measured by the Gini coefficient, a growing maximum node degree, a slightly falling average local clustering coefficient, a nearly constant small diameter of the largest connected component, and a slightly growing number of connected components (which can be explained by some small connected components that exist in addition to a giant component). We detect communities using PLM (Parallel Louvain Method), a modularity-based community detection heuristic [65], and report the number of communities minus the number of these small connected components. It can be observed that the number of non-trivial communities grows slightly.

While we do not propose that these scaling laws are universal, the trends represented here are commonly observed [15, 9, 62]. Thus, we use them to define desired scaling properties for the remainder of the study as follows: m grows linearly with n; the diameter does not change significantly, preserving the "small world property"; the shape of the degree distribution remains skewed; the maximum node degree increases; the number of connected components may grow; the number of communities increases slightly.

Recall that one use case for our generator is testing of graph and network analysis algorithms. Since the running time is an essential feature in such tests, we also consider a realistic replication of running times important. To this end, we select a set of graph algorithms that (i) compute important features of networks and are thus frequently used in network analysis tasks and that (ii) cover a variety of patterns of computation and data access, each of which may interact differently with the graph structure. The set consists of algorithms for connected components (essentially breadth-first search), PageRank (via power iteration), betweenness approximation (according to Geisberger et al. [25]), community detection (PLM, [65]), core decomposition (according to [20]),

triangle counting (according to [30]), and spanning forest (essentially Kruskal's algorithm without edge weights).

3 The Generation Algorithm ReCoN

We introduce ReCoN, a generator for replicating and scaling complex networks. Its input is a graph and a community structure on it. For fitting a given graph without given community structure, we use PLM [65] in order to detect a community structure first. The basic idea of ReCoN is to randomize the edges inside communities and the edges between communities while keeping the node degrees. This happens separately such that each community keeps as many edges as it had before. For scaling a graph, we first create as many disjoint copies of the graph as desired and then apply the aforementioned steps. During the randomization of the edges between the communities the copies usually become connected with each other.

The idea of randomizing graphs inside and between communities is inspired by the LFR generator, a benchmark graph generator for community detection algorithms [38]. There the basic building blocks are also a random subgraph per community and a global graph. However, in the LFR generator the degrees and communities are not given but generated using a power law degree distribution and a power law community size distribution with nodes assigned to communities at random, while ReCoN uses the given graph as input for them.

For randomizing graphs while preserving the degree sequence, we use random edge switches where two edges $\{u, v\}$, $\{y, z\}$ chosen uniformly at random are changed into $\{u, z\}$, $\{y, v\}$ if the resulting graph is still simple, i. e. does not contain any duplicate edges or self-loops. Similarly to the edge switching implementation provided by [70] we use 10 times the number of edges as the number of random edge switches. Previously performed experiments (e. g. [51]) have shown that this is enough to expect the resulting graph to be drawn uniformly at random from all graphs with the given degree sequence.

For an original graph $O = (V, E)$ with $n_o = |V|$ nodes and a desired scaling factor x, ReCoN executes the following steps:

1. Detect a community structure $\mathcal{C} = \{C_1, \ldots, C_k\}$ on O using PLM.
2. Create H as the disjoint union of x copies of O. The community structure is also copied such that the new community structure $\mathcal{D} = \{D_1, \ldots, D_{x \cdot k}\}$ consists of $x \cdot k$ communities, i. e. each copy of O gets its own copy of the community structure that is aligned with the structure of the copied graph.
3. For each community D_i, $1 \leq i \leq x \cdot k$, randomize the edges of the subgraph $H[D_i]$ that is induced by the community D_i while keeping the degree distribution using random edge switches.
4. Randomize the remaining edges, i. e. all edges in H that are not part of one of the subgraphs $H[D_i]$ using random edge switches. Note that afterwards some edges that were not in one of the $H[D_i]$ can now be inside a community. In order to avoid this, rewiring steps are performed by executing edge switches of such forbidden edges with random partners. A similar step is also used in the LFR generator, where it was observed that in practice only few rewiring steps are necessary [37].

Note that it is not necessary to start with the full information of the original graph in Steps 3 and 4. It is actually sufficient to know a community structure (as opposed to the whole original graph) and for each node the internal and external degree, i. e. how many neighbors it has inside and outside its community, respectively. For our implementation we choose this alternative. Further, we execute Step 3 in parallel for all communities as the subgraphs are disjoint.

In addition to replicating important properties with high fidelity, the randomization in Steps 3 and 4 naturally produces random variance among the set of replicas.

3.1 Fitting Generative Models to Input Graphs

Parametrized generative models require learning model input parameters from the original network. A fitting scheme is an algorithm that takes a network as input and estimates parameters of a generative model. Because, usually, such schemes are not unique, exploring them would be important future work. For this study, we have chosen one straightforward scheme per model, parameters of which are summarized in Table 1. Below we discuss a fitting scheme for power law degree distributions, and briefly describe the generative models that are compared with ReCoN.

We consider an original graph $O = (V, E)$ with $n_o = |V|$, $m_o = |E|$ and maximum degree $d_{\max}$. We denote as $(a_i)_{i \in M}$ a sequence of elements a with indexes i from an (ordered) set M, so

$$\bigcup_{j=0}^{k} (a_i^j)_{i \in M}$$

denotes the ordered concatenation of k sequences. The clustering coefficient of node u is denoted by c_u. The number of nodes with degree d is denoted by

$$n_d = |V_d| = |\{v \in V : \deg(v) = d\}|,$$

and

$$c_d = \frac{1}{n_d} \sum_{u \in V_d} \mathsf{c}_u$$

corresponds to the average clustering coefficient for nodes of degree d.

After a brief description of the considered models, we discuss fitting of power-law distributions, and continue with a discussion of the parametrization of each model.

Model	**Parameters**	**Fitting**	**Fitting** Scaling by $x \in \mathbb{N}$
Erdős–Rényi	$ER(n', p)$	$n' = n \qquad p = \frac{2m}{n \cdot (n-1)}$	$n' = x \cdot n \qquad p = \frac{2m}{x \cdot n \cdot (n-1)}$
Barabasi-Albert	$BA(n', k)$	$n' = n \qquad k = \lfloor m/n \rfloor$	$n' = x \cdot n \qquad k = \lfloor m/n \rfloor$
Chung-Lu	$CL(d)$	$d = (\deg(u))_{u \in V}$	$d = \cup_{i=1}^{x} (\deg(u))_{u \in V}$
Edge-Switching Markov Chain	$EMC(d)$	$d = (\deg(u))_{u \in V}$	$d = \cup_{i=1}^{x} (\deg(u))_{u \in V}$
RMAT	$RM(s, e, (a, b, c, d))$	$s = \lceil \log_2 n \rceil \qquad e = \lfloor m/n \rfloor$ $(a, b, c, d) = \mathsf{kronfit}(O)$	$s = \lceil \log_2 x \cdot n \rceil \qquad e = \lfloor m/n \rfloor$ $(a, b, c, d) = \mathsf{kronfit}(O)$
Hyperbolic Unit-Disk	$HUD(n, \bar{d}, \gamma)$	$n = n \qquad \bar{d} = 2 \cdot (m/n)$ $\gamma = \max\{2.1, \mathsf{plfit}((\deg(u))_{u \in V})\}$	$n = x \cdot n \qquad \bar{d} = 2 \cdot (m/n)$ $\gamma = \max\{2.1, \mathsf{plfit}((\deg(u))_{u \in V})\}$
BTER	$BTER(d, c)$	$d = (n_d)_{d \in (0, \ldots, d_{\max})}$ $c = (c_d)_{d \in (0, \ldots, d_{\max})}$	$d = (n_d \cdot x)_{d \in (0, \ldots, d_{\max})}$ $c = (c_d)_{d \in (0, \ldots, d_{\max})}$
LFR	$LFR(n', \gamma, \bar{d}, d_{\max}, \beta, c_{\min}, c_{\max})$	$n' = n$ $\gamma, d_{\min}, d_{\max} = \mathsf{plfit}((\deg(u))_{u \in V})$ $\zeta_s = \{\lvert C \rvert \mid C \in \mathsf{PLM}(O)\}$ $\beta, c_{\min}, c_{\max} = \mathsf{plfit}^*(\zeta_s)$	$n' = x \cdot n$ $\gamma, d_{\min}, d_{\max} = \mathsf{plfit}((\deg(u))_{u \in V})$ $\zeta_s = \{\lvert C \rvert \mid C \in \mathsf{PLM}(O)\}$ $\beta, c_{\min}, c_{\max} = \mathsf{plfit}^*(\zeta_s)$

Table 1: Parameters set to fit a model to a given graph, and to produce a scaled-up replica.

3.2 Considered Models

Apart from ReCoN, we include the following models in our experimental study: *Erdős–Rényi* random graphs (ER) [22], *Barabasi–Albert* (BA) [3], *Chung-Lu* (CL) [1], *Edge-Switching Markov Chain* (ESMC), *Recursive Matrix* (RMAT) [18], *Hyperbolic Unit-Disk Graph* (HUDG) [35], BTER [61], and *LFR* [38].

3.2.1 Fitting power law degree distribution (PLD)

Both LFR and HUDG generators produce graphs with a power law degree distribution. Therefore, at least the power law exponent, and, in the case of the LFR generator, also the minimum and maximum degrees need to be determined such that the degree distribution fits the real network. In Table 1, plfit refers to our custom power law fitting scheme.

A practical replication of a network requires preserving the original average (otherwise, the density will be changed) as well as minimum and maximum degrees (applications can be sensitive to such fundamental properties as degree-1 nodes and the distribution of hubs).

In general, it is assumed (and implemented in many algorithms [19]) that PLD only holds from a certain degree on and that for smaller degrees, the distribution might be different. As the LFR generator only generates a plain PLD, we cannot apply this assumption. Therefore, we fit the PLD exponent such that, with the given minimum and maximum degree, the average degree of the real network is expected when a degree sequence is sampled from this PLD. Using binary search in the range of $[-6, -1]$ we repeatedly calculate the expected average degree until the power law exponent is accurate up to an error of 10^{-3}.

3.2.2 Erdős–Rényi, Barabasi-Albert, Chung-Lu and ESMC

Erdős–Rényi random graphs (ER) [22], while not considered similar to most real-world graphs, are fundamental and an important baseline.The Erdős–Rényi model does not provide many options for parametrization. The edge probability p is set to produce the same edge-to-node ratio $\frac{m}{n}$ as the original. The *Barabasi–Albert* model (BA) [3] implements a preferential attachment process by which a PLD emerges, which has been claimed to be a typical feature of real complex networks. In BA, we set the number of edges coming with each new node to fit the original edge-to-node ratio. The *Chung-Lu* (CL) model [1] recreates a given degree sequence in expectation. The *Edge-Switching Markov Chain Generator* (ESMC) generates a graph that is randomly drawn from all graphs with exactly the given degree sequence (see e.g. [51], [60]). In both CL and ESMC we use the original degree sequence. To generate larger networks, x copies of this sequence are concatenated, multiplying the number of nodes by x while keeping the relative frequency of each degree.

3.2.3 RMAT

The *Recursive Matrix* (RMAT) model [18] was proposed to recreate various properties of complex networks, including an optional power-law degree distribution, the small-world property and self-similarity. The RMAT model can only generate graphs with 2^s nodes, where s is an integer scaling parameter. In order to target a fixed number of nodes n_r, we calculate s so that $2^s > n_r$ and delete $2^s - n_r$ random nodes. The choice of the parameters a, b, c, d requires some discussion.

Leskovec et al. [41] propose a method to "given a large, real graph [...], generate a synthetic graph that matches its properties", using stochastic Kronecker graphs: Starting 2-by-2 stochastic initiator matrix I, Kronecker products are calculated so that I^s is a stochastic matrix of dimension 2^s that yields edge probabilities of a graph. This is equivalent to the RMAT model as it yields the same edge probabilities. They attempt to fit model parameters so that the likelihood that a given original graph O was generated starting from an initiator matrix I is maximized, and propose the kronfit gradient descent algorithm that iteratively estimates I in $O(m)$ time. They do not explicitly mention the case of creating a scaled replica, but it is clear that the method is capable of producing graphs for arbitrary exponents s. We use an implementation of kronfit which is distributed with the SNAP network analysis tool suite [43].

The high running times of kronfit (see Figure 2) motivate a closer look at the claim of increased realism as opposed to random guessing of the initiator matrix. We compared properties of original networks, RMAT replicas with either random initator matrices or matrices generated with 50 iterations of kronfit (a value proposed in [41]), see Appendix A. The most obvious difference between original and replicas is in the distribution of the local clustering coefficient, which is consistently lower for RMAT. Applying kronfit diminishes the mismatch slightly, but not fundamentally. This points to inherent restrictions of generating graphs by stochastic Kronecker multiplication/RMAT

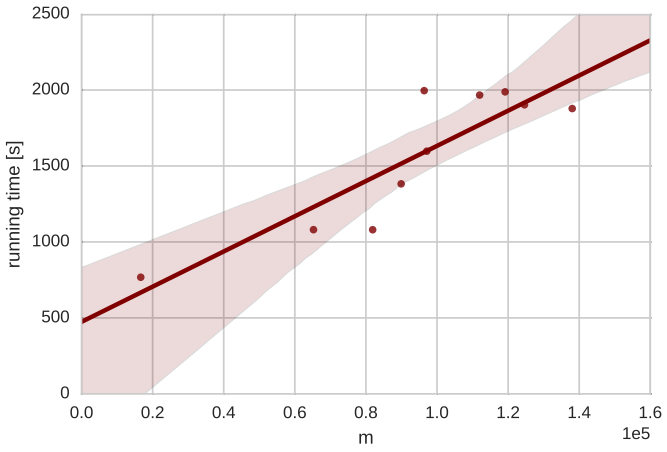

Figure 2: Running time of 50 iterations of the kronfit algorithm in relation to the number of edges m of the input network

that an elaborate parameter fitting does not overcome. Since running kronfit on every network to be replicated is not practical, we estimate RMAT parameters as follows: We assume that the 100 Facebook networks constitute a class with essential structural commonalities, and run the recommended 50 iterations of kronfit on one typical network, `fb-Caltech36`. The resulting initiator is applied to replicate all Facebook networks. For other sets of networks, the assumption of structural similarity cannot be made, so we use a new random initiator for each replication.

3.2.4 Hyperbolic Unit Disk Graphs (HUDG)

The random hyperbolic graph model embeds nodes into hyperbolic geometry and connects close nodes with higher probability [35]. The unit-disk variant HUDG we use in this paper connects only nodes whose distance is below a certain threshold. We are focusing on the unit-disk variant to be able to use a very fast generator for this model [71]. The model has been shown to replicate some properties observed in real networks, such as a power-law degree distribution. This method receives as parameters the desired number of nodes, the average degree of the original network and a power law exponent which is fitted as described above. As the given power law exponent must be larger than 2, we supply at least an exponent of 2.1.

3.2.5 BTER

The BTER model [61] receives a degree distribution and the desired clustering coefficient per degree, i.e., for each degree to be realized the number of occurrences and the average clustering coefficient per degree. For scaled replicas we scale the occurrences of all degrees by the scaling factor. This leads to the target number of nodes while also preserving the general shape of the degree distribution. In order to retain the distribution of the clustering coefficients, we leave them unchanged while scaling the network.

3.2.6 LFR

The LFR was designed as a benchmark graph generator for community detection algorithms [38]. It generates graphs with a power law degree distribution and a community structure with community sizes that follow a power law distribution. Apart from the number of nodes it requires parameters for power law distributions of the node degrees and the community sizes, and a mixing parameter that determines the ratio between intra- and inter-cluster edges. We detect communities using

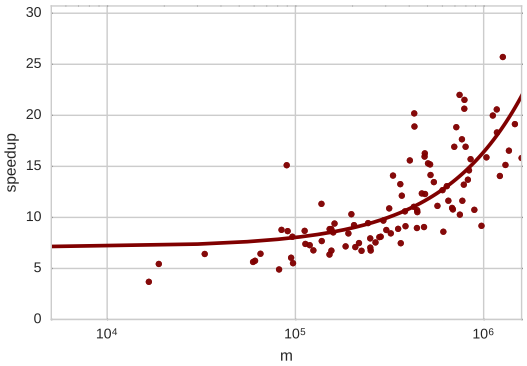

Figure 3: Speedup of NetworKit implementation of LFR compared to the reference implementation [24] when replicating a set of networks with m edges

PLM [65] and fit the parameters for the two power law distributions as described above using the original degree sequence and the found community sizes. The mixing parameter is set to the ratio between intra- and inter-cluster edges of the found communities. As many networks in our set of real-world networks contain a few small connected components, the smallest communities usually only contain just 2 nodes while all other communities are much larger. Therefore, frequently our minimum power law exponent 1 is chosen as exponent and the expected average community size is still too low. In these cases we also fit the minimum community size using binary search until the expected average community size is the real average community size. In Table 1 this is denoted by plfit*. For the scaling, we simply append copies of the same values.

4 Computational Experiments

Our implementations of ReCoN and the various fitting methods are based on NetworKit [66], a tool suite for scalable network analysis. It also contains many of the generators we use for comparison and provides a large set of graph algorithms we use for our experiments. NetworKit combines C++ kernels with an interactive Python shell to achieve both high performance and interactivity, a concept we use for our implementations as well. All implementations are freely available as part of the package at `https://networkit.iti.kit.edu`. This also includes a faster and parallel implementation of the LFR generator (compared to the original implementation [24]).

Our experimental platform is a shared-memory server with 256 GB RAM and 2x8 Intel(R) Xeon(R) E5-2680 cores at 2.7 GHz, using the GCC 4.8 compiler and the openSUSE 13.1 OS. As described in Section 2, we are interested in how well the different generators replicate certain structural features of the original networks as well as the running times of various graph algorithms. The results are described subsequently.

4.1 Implementation details

A reference implementation of the LFR generator by Fortunato et al. is available online [24]. We created a custom implementation of LFR in NetworKit, which is also the first parallelized implementation. Fig. 3 illustrates the speedup that we achieve with a new implementation of the LFR generator based on NetworKit. Speed measurements were obtained on a machine with 16 physical cores. It yields speedup factors of 5 to 25 for the test set of 100 Facebook social networks, and the factor grows superlinearly with the size of the network to be generated. The running time difference can be partially traced back to implementation differences: The LFR reference implementation relies on `std::set` to store and test for graph adjacencies, while our implementation uses a simple but apparently more efficient sequential scan on a simple array. Further, as mentioned already, the community graphs are generated in parallel, which gives an additional speedup.

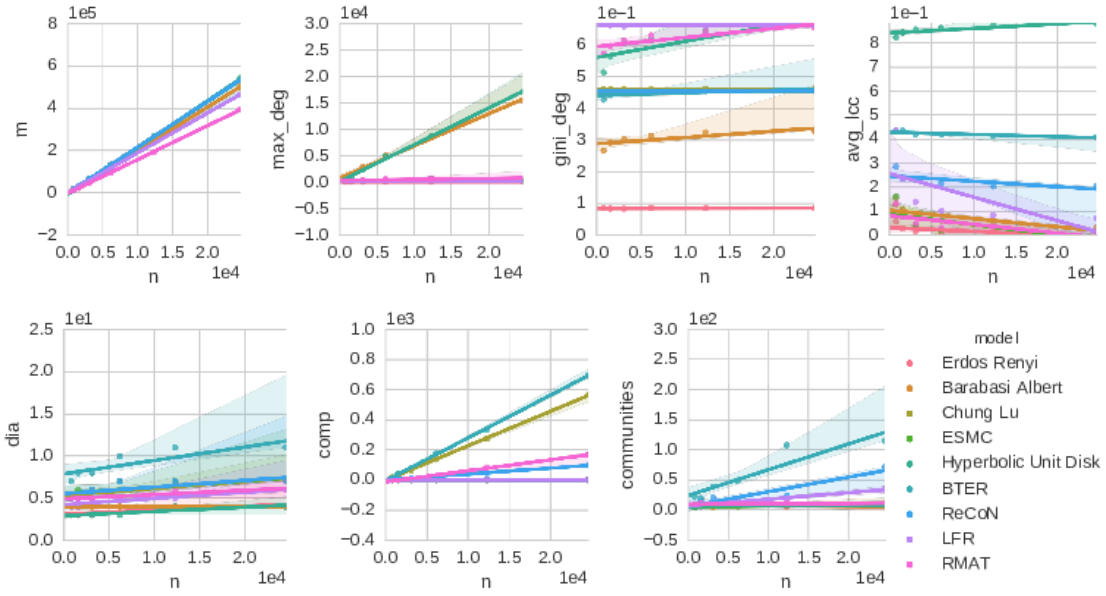

Figure 4: Scaling behavior of the different generators on the `fb-Caltech36` network. From left to right and top to bottom: number of edges, max. degree, Gini coefficient of the degree distribution, average local clustering coefficient, diameter, number of components, number of communities.

For the tested generative models, NetworKit includes efficient implementations, such as implementations of linear-time algorithms for the Erdős–Rényi and Barabasi-Albert models [7] and a subquadratic time algorithm for HUDG [71]. We have not implemented the BTER model; instead we use the FEASTPACK implementation by Kolda et al. [67]. NetworKit implements an adapter class that performs the model fitting and transparently calls the MATLAB-based FEASTPACK binary.

4.2 Scaling behavior of the generators

The following experiments consider the scaling behavior of generative models. Given the parametrization discussed before, we look at the evolution of structural features with growing scale factor x up to $x = 32$. We consider the same basic scalar features as for the real networks in Section 2.

In Figure 4, we show the results of the scaling experiments for the `fb-Caltech36` network. The number of edges of the replicas is increased almost linearly by all generators to $\approx 5 \cdot 10^5$ edges, which approximately corresponds to 32 times the edges of the original network. Therefore, all generators seem to keep the average degree of the original network, which is expected as it is a parameter of all considered generators. Surprisingly, the maximum degree strongly increases up to 10 or 15 thousand with HUDG and BA generators, respectively. The original maximum degree is 248, so that the new value is even significantly higher than the scaled maximum degree (i. e. $248 \cdot 32$). Actually, from the scaling study in Sec. 2, we could expect an increase, but rather in a lower range, so the degree distribution of BA and HUDG generators are not realistic. Concerning the Gini coefficient, one can clearly see that ER does not generate a skewed degree distribution at all. All generators that get the exact degree sequence as input keep the Gini coefficient constant, which is expected and also relatively realistic from our scaling study.

The original average local clustering coefficient of 0.43 is almost exactly reproduced by BTER in which it is an input parameter. The HUDG method increases it to 0.8, most others obtain very small values. Our new ReCoN generator is less far off with 0.25 and a slightly decreasing clustering coefficient; the latter is actually realistic as we saw in Sec. 2. LFR is able to generate a clustering coefficient above 0.2 initially. Other generators produce much lower clustering coefficients. The original diameter of 6 is almost exactly kept by ReCoN, all other generators except BTER produce networks with slightly lower diameters, while BTER generates networks whose diameter is almost

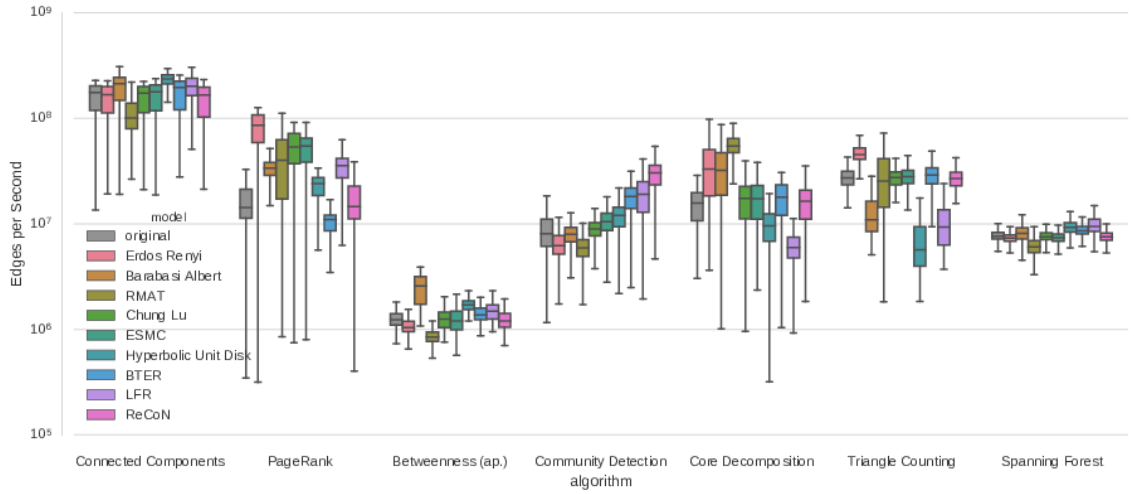

Figure 5: Running time replication of a set of network analysis algorithms. Running times are in edges per second, i.e., higher is faster.

twice bigger. All generators show a slight increase of the diameter when the networks are larger, which is consistent with our scaling study. While most generators produce networks with just a single connected component, CL and BTER generate a large number, RMAT and ReCoN a moderate number of connected components. In the case of CL, BTER and RMAT, this is probably due to a large number of degree-0 nodes. The original network consists of a giant component and 3 small components; ReCoN scales them linearly, which is due to its parametrization. The original network is split into eight non-trivial communities, that number should increase slowly according to Sec. 2. Only in the networks generated by BTER, ReCoN and LFR, PLM can find a significant and increasing amount of communities. While PLM finds over 100 non-trivial communities in the network generated by BTER, there are fewer communities detectable in the networks generated by ReCoN and even less in the ones generated by LFR. Overall, ReCoN is the only generator that keeps the degree distribution, and produces a realistic clustering coefficient and a small diameter while keeping the graph connected and preserving a moderate number of communities. All other generators are either unable to keep the diameter or the connectivity or the number of communities. It is part of future work to investigate whether the full hyperbolic random graph model can alleviate the weaknesses of the unit-disk case.

In Appendix B we present results of further experiments on the preservation of different network properties both for scale-1 and for scale-4 replicas. The conclusion there is similar: for some properties, ReCoN is outperformed by other generators but overall it is able to produce the most similar replicas.

4.3 Realism in the running times of graph algorithms

Synthetic graphs are frequently used in algorithm engineering to estimate the running time of an algorithm assuming that this time will be similar on real networks. We examine if this is indeed the case with the generative models we consider. Using the previously described generators and fitting schemes, we generate replicas of 100 Facebook networks and test a variety of graph algorithms (see Sec. 2) on both the original and replica sets.

Our experiments demonstrate (see Fig. 5) that the running times on the replica sets often do not match those on the original set. The gray segments of the box plots represent the distribution of running times measured on a set of original networks. Ideally, the distribution on the synthetic networks would be identical. The difference is statistically nontrivial, though. Small variance between the models exists for connected components and spanning forest computations, since their running time is nearly constant per edge. Other algorithms exemplify how much running time can depend on network structure, especially community detection, core decomposition, triangle counting and PageRank. In general, the running time measurements obtained on ReCoN match the originals closely in most cases. An exception is community detection, where PLM seems to

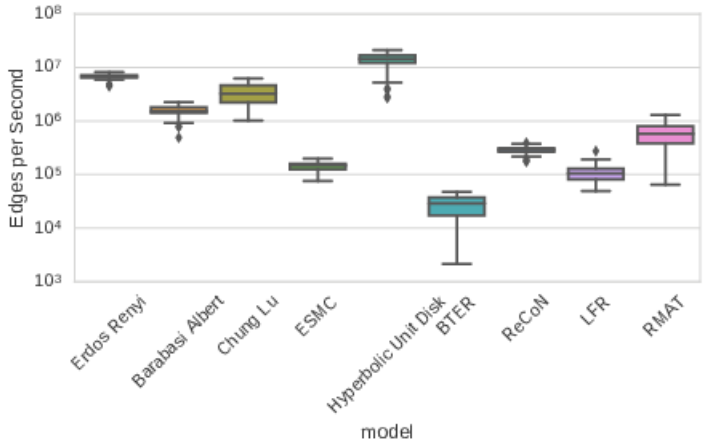

Figure 6: Fitting and generating: processing speed measured in edges per second (size of replica graph/total running time, measured on 100 Facebook graphs)

profit from ReCoN's explicit model of communities. BTER shows close matches, too.

4.4 Generator running times

In Fig. 6, we show the running times of parameter fitting and generating a replica for all methods. Processing speed is given in the number of edges per second, a common measure for benchmarking graph algorithms that sets graph size and execution time in relation. The entire set of Facebook networks was used to produce the measurements, so generated replicas range from about 15000 to 1.5 million edges. For all models, generating the graph takes up the vast majority of time. BTER's MATLAB-based implementation is slowest, while the ER and HUDG generators are the fastest. Our implementations of LFR and ReCoN are not among the fastest generators, but fast enough to produce millions of edges in minutes.

4.5 Social Network: Example Replication I

As part of the evaluation of realism we replicate a small example network with each of the models. The input graph is a social network of Bottlenose dolphins. We use graph drawing with the ForceAtlas2 layout algorithm implemented in Gephi to inspect the network structure. Node sizes are proportional to degree within one drawing. Fig. 7(a) shows a layout of the original network. Consider its basic structure: The graph is connected, with two distinct communities separated by an edge bottleneck. We also layouted replicas with scaling factors of 1 and 2 for each model, finding that the ReCoN replicas (e.g. Fig. 7(d)) are the only ones commonly recognized as resembling the original. BTER (Fig. 7(c)) is able to generate a similar community structure, though not matching the original as closely. All other replicas (e.g. those created by RMAT after applying kronfit, Fig. 7(b)) lose the distinctive community structure of the original.

4.6 Epidemiological Network: Example Replication II

An epidemiological network frequently used in studies on the transmission of HIV is based on data collected in Colorado Springs by Potterat et al. [58]. It contains 250 individuals who were in contact in the 1980s through sex or injection drug use. Figure 8(a) shows a force-directed layout of the network's graph. Characteristic for this network is its tree-like structure and the presence of high-degree hubs with attached "satellite" nodes of degree 1. The Colorado Springs network is an instance of network data that cannot be shared freely due to legal restrictions, but no such restrictions apply to a synthetic replica. A replica made by the ReCoN generator reproduces structural features of the original with the highest accuracy among the considered

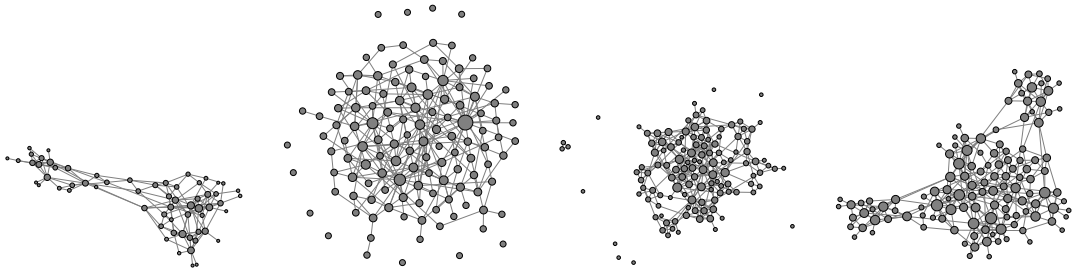

(a) original (b) RMAT with kronfit (c) BTER (d) ReCoN

Figure 7: A small social network and its scale-2 replicas produced by different models. ReCoN is the model that best reproduces a set of essential properties, including degree distributions, clustering and community structure.

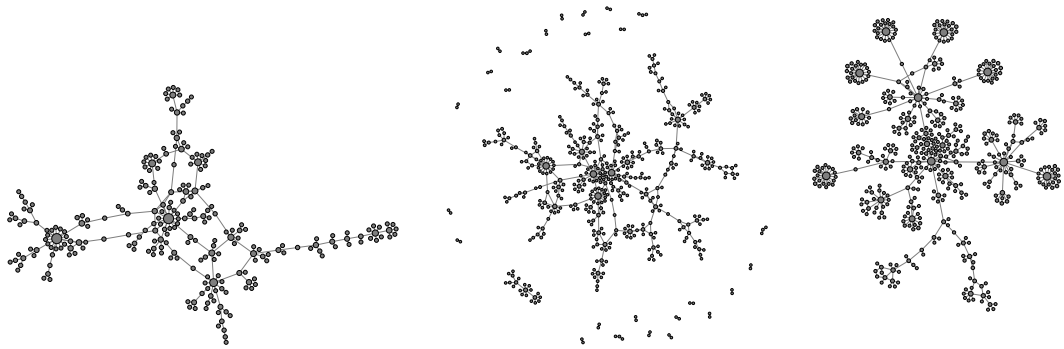

(a) original (b) ReCoN, scale-2 (c) ReCoN, scale-200k (sample)

Figure 8: Colorado Springs epidemiological contact network

models. More importantly, scaled replicas retain these essential properties, including the hub-and-satellite structure. Figure 8(b) shows the network replicated with a scaling factor of 2, which is remarkably close to the original. Like several other generators, ReCoN generates additional small connected components as an artifact. If this is undesirable, a postprocessing step could prune the network down to its giant connected component containing a large majority of nodes. (Interestingly, the original data set contained many singletons and isolated dyads that were removed from the graph in a preprocessing step).

Real epidemiological contact network data is difficult to collect, further complicated in the case of HIV by sex and drugs being tabu subjects. This makes obtaining such a network on the scale of an entire population impractical. In such a scenario, the ability to create realistic large synthetic replicas of smaller real networks may be highly relevant. As an explorative case study, we let ReCoN generate a replica of the Colorado Springs network with 50 million nodes, which corresponds to a scaling factor of 200 000. Figure 8(c) shows a sample from a $5 \cdot 10^7$ million node replica. This network's structure is quite different from the clustered "friendship networks" of dolphins and humans (Facebook). Remarkably, ReCoN also replicates many aspects of the original's structure very closely, such as a tree-like structure with hubs and attached satellites. These aspects are retained even for huge scaling factors. This makes ReCoN a promising candidate to deliver large data sets in cases where large amounts of real data are likely unobtainable. Further domain-specific validation of the suitability of such replicas would be interesting.

5 Multi-scale modeling and open problems

The ReCoN algorithm introduced above belongs to the class of single-scale generators in which the network structure is replicated in breadth in order to achieve an increased synthetic network without considering structural properties at coarse-grained resolutions of the network. That is, the community structure is replicated multiple times, and connections are randomized while preserving key properties of the original network. The approach allows us to extrapolate from a given network to the next level up; however, it is clear that as a network grows, its structure evolves new coarse-grained resolutions (or hierarchical levels) on top of those that were already present.

Progress in modeling such multiscale networks has been reported in [29], which introduced a multiscale network generation approach, MUSKETEER. Briefly, starting from a single known or hypothesized network from any domain, MUSKETEER synthesizes ensembles of networks that preserve, on average, a diverse set of structural features *at multiple coarse-grained scales*, including degree, other measures of centrality, degree assortativity, path lengths, clustering, and modularity, while introducing unbiased variability across the ensemble in many of these properties. The core method is inspired by applications of the theory of multiscale methods [10, 13, 69] to combinatorial optimization [11, 47, 59]. The original computational multiscale method starts with an optimization problem represented on a network, constructs a hierarchy of decreasing in size networks $G_0 = O, G_1, ..., G_k$ via a coarsening procedure [59], solves a small optimization problem at the coarsest scale, and then iteratively uncoarsens the solution by gradual prolongations and local refinements from the coarse to next-finer scales. Similarly, in MUSKETEER a hierarchy of coarse networks is created; but, in contrast to the multiscale methods for computational optimization, nothing is optimized but the network is edited at all scales of coarseness. During the editing process only local changes are allowed, which are the results of local decisions only. In other words, the problem of network editing/replication/randomization is formulated and solved at all scales, where primitives at the coarse scale (such as coarse nodes and edges) represent aggregates of primitives at previous finer scale. Analogous to multiscale methods for computational optimization problems [10, 11], by using appropriate coarsening, we are able to detect and use the "geometry" behind the original network at multiple scales, which can be interpreted as an additional property that is not captured by other network generation methods customized to replicate a small number of well known properties of O such as degree distribution. Moreover, it is known that the topology of many complex networks is hierarchical and thus might be produced through iterations of generative laws at multiple scales. In general, such generative laws often can be different at different scales, as evidenced by the finding that complex networks are self-dissimilar across scales [16, 31, 56, 73, 8, 52]. These differences can naturally be reflected in the multiscale frameworks.

We believe that multiscale network modeling methods (MNM) are particularly promising for addressing open problems in MNM, and we outline them below. Some of the largest open problems are in the areas of statistical network hypothesis testing, network anonymization and compression, multiscale network decomposition and comparison, and understanding the fundamental limitations of network modeling.

5.1 Statistical network hypothesis testing

One of the biggest unsolved problems in network science is the problem of hypothesis testing. This is a problem of high significance in fields such as infectious diseases as well as other research and its application. Given a network, the problem is to determine if the given network has statistically significant differences with a background network. For example, given a security problem on a network, could we say that a given network is significantly more "robust" (or any macroscopic property) than what would be expected by chance, where "chance" is defined realistically? The only existing approach, which is unsuitable for most applications, is to fit a parametric model such as an ERGM to the network, and then test for the parameters being different from 0. This ERGM approach cannot be used in the vast majority of applications, because the network cannot usually be represented realistically using an ERGM, nor could most properties be represented easily using a generative parameter.

However, given the high realism and versatility achieved by MNM methods, we believe that MNM, and non-parametric models more generally, could significantly advance statistical network

hypothesis testing. The non-parametric method promises to provide a much more realistic permutational test, and therefore identify anomalous structures in a given network at a much lower level of false positives.

5.2 Network anonymization and compression

A critical barrier towards advances in network science is the difficulty of sharing network data, which is often confidential and sensitive. This is the problem of network anonymization and compression. The multiscale strategy promises to address this gap through a powerful non-parametric data synthesis strategy, because, as we previously demonstrated, the strategy can generate networks that can model network structures with a high level of realism across a large spectral range. Initial success with such methods is possible already with the ReCoN and MUSKETEER strategies, but much work remains to be done addressing the spectrum of anonymization needs and developing versatile algorithms. Current algorithms operate with complete network exemplars or replicas, but it could be beneficial for anonymization to operate with only small-size descriptions of the network such as motifs or possibly global parameters.

The MNM-based anonymization could be either non-parametric or parametric, allowing the user to limit the amount of data released to the public. In the parametric version, the data owner shares a controlled amount of network information, while the recipient is able to obtain a realistic representation of the data by using an MNM algorithm.

Such a system would offer major advances in areas such as public health, cybersecurity and counter-terrorism. Data in these areas is often very difficult to collect and/or locked by privacy laws. Anonymization would compress the sensitive datasets and generate non-sensitive data, and in this form they could be shared with researchers in multiple institutions and even the public. This would vastly expand the number of people who can perform research in critical areas of scientific and practical interest.

5.3 Multiscale decomposition and comparison of networks arising from real-world complex systems

The complexity of real network data suggests that its structure would be generally self-dissimilar when comparing several scales (or coarse-grained resolutions). However, existing network generators generally fail to address the problem of self-dissimilarity. For example, the degree distribution or the clustering coefficient of a communication network are not necessarily the same as those at its backbone at macro- and meso-scopic scales. Moreover, in the vast majority of practical cases, the coarse-grained structure is not given (or can be found) explicitly, which makes the attempts to generate realistic synthetic network even more complicated. To address this problem, some of us and others have investigated a variety of multilevel optimization solvers for such problems, including graph partitioning [14, 26, 48], linear ordering [59], graph drawing [47], and node immunization [44] in which the multiscale organization of a graph is expressed via low-energy cuts. However, other advanced approaches are needed when a decomposition into scales is not cut-based.

A related open problem has to do with the comparison of networks with self-dissimilar properties. In the existing multiscale algorithms like ReCoN and MUSKETEER, the comparison of the original and synthetic networks is done at the finest scale only. Whereas in some applications this can be acceptable, clearly, there are areas in which the multiscale structure is considered, and, thus, a multiscale comparison is much more beneficial and illuminating. However, to the best of our knowledge, there is no comprehensive concept of similarity that takes into account a comparison of properties at different scales.

5.4 Fundamental limitation of network modeling algorithms

Despite many studies of network modeling, we appear to lack a fundamental theory to understand the limitations of these algorithms. There are at least three related problems. First, we lack a theoretical understanding for many of the most popular generating algorithms, including multiscale strategies such as MUSKETEER. Second, we lack a sense in which we can describe a network model or generator as in some sense "optimal" or "most realistic". Perhaps it is possible to give an

information-theoretic basis to understand the maximal realism a given generator can achieve given its parameter space. Third, complex networks are typically associated with a variety of different processes running on or creating them. Nearly all network generation models focus on rules that are structuring the network, overlooking the close dependence between the structure-forming processes and the processes on the network. Indeed, in many domains, the realism of a network is closely tied to the dynamics on the network. A domain expert (i.e., a potential user of network generation methods) cannot consider a synthetic graph realistic unless it becomes obvious that the processes that are supposed to co-exist with the network are also realistic (an issue we examined in the context of ReCoN). It is hoped that future studies in this field would examine this and earlier gaps in detail.

6 Conclusion

We have presented a new generator, ReCoN, for replicating and scaling existing networks. In an extensive experimental evaluation we have shown that ReCoN is capable of generating networks which are (i) similar to the original network in terms of important structural measures and (ii) lead to similar running times of many graph and network analysis algorithms. Using ReCoN it is possible to realistically replicate an existing network, and to scale the synthetic version by orders of magnitude, e. g., in order to test algorithms on larger data sets where they are not available. Furthermore, it allows to create anonymized copies of such networks that can be distributed freely and allow to conduct representative experiments on them. While other generators sometimes perform better concerning certain criteria, none of the other generators is capable of approximately reproducing such a wide range of properties and running times.

A Appendix: Structure profiles for the evaluation of the kronfit scheme

The high computational cost of kronfit motivates a closer look at the claim of increased realism as opposed to random guessing of the initiator matrix. Network profiles generated by NetworKit enable a quick exploratory analysis of an original network and replicas created by RMAT, with random parameters on the one hand and parameters generated after 50 iterations of kronfit (a value that the authors propose in [41]). The original is `fb-Caltech36`, a small instance from the "Facebook 100" collection of social networks. Figures 9, and 10 show an excerpt from the generated profiles. The most obvious difference between the original and its two replicas is in the distribution of the local clustering coefficient: While the median is around 0.4 for the original, it is about 0.1 for the replica with random parameters and only slightly higher for the konfit-generated replica. The same behavior of the generator was observed consistently on other networks. This points to inherent restrictions of generating graphs by stochastic Kronecker multiplication/RMAT that an elaborate parameter fitting does not overcome.

Since running kronfit on every network to be replicated is not practical, we estimate RMAT parameters as follows: We assume that the 100 Facebook networks constitute a class with essential structural commonalities, and run the recommended 50 iterations of kronfit on one typical network, `fb-Caltech36`. The resulting initiator

$$I_{\texttt{fb-Caltech36}} = \begin{bmatrix} 0.378802757 & 0.249474498 \\ 0.255098510 & 0.116624233 \end{bmatrix}$$

is applied to replicate all Facebook networks. For other sets of networks, the assumption of structural similarity cannot be made, so we use a new random initiator for each replication.

B Appendix: Replicating Structural Properties

In the existing literature on generative models, claims of realism are typically substantiated by showing that a set of structural properties is similar for real and synthetic networks. The large palette of properties to choose from and the question which of those properties are essential features makes this a complex problem. An often used approach is to describe a network by a feature vector, a set of scalar properties. These are often maxima, minima or averages of node properties (cf. the structural profiles published by the KONECT project [36]). This can be reductionist, since these summary values may not give enough information about how the node properties are distributed. We therefore demonstrate how well the different models replicate networks with two types of plots. The first plot type shows scalar properties of the replicas, relative to those of the original. Let x_o and x_r denote scalar properties of the original and the replica, respectively. The relative value x_r/x_o is computed for each replica. Box plots depict the distribution of these relative values over the entire set of replicas. Scalar properties included are: the number of edges, the number of connected components, the effective diameter (for 90% of node pairs) of the largest connected component, and the number of communities. The effective diameter is approximated using the ANF algorithm [57]. Communities are detected using the modularity-maximizing algorithm PLM [65]. This set was chosen so that it could be quickly computed for a large set of networks. The second type of plot covers centrality measures, and is designed to show how the shape of the distributions of node centrality scores of the originals compares to those of the replicas. Each segment of the plot depicts the centrality values of all nodes of all networks in the considered data set or the replicas of a certain algorithm. Since centrality measures from this selection can have very different scales, all centrality scores are normalized to the interval $[0, 1]$.

Describing the results on 100 Facebook networks (Fig. 11(a)) from left to right, we observe the following results for the measured scalar properties: All models were parametrized to produce exactly $n' = n$ nodes, while some degree of freedom exists in the number of edges. However, all generators replicate the number of edges with a narrow variance, which is largest for RMAT. ReCoN is the only model that matches the number of components with high accuracy, while it is lower for ER, BA, ESMC, HUDG and LFR, and significantly higher for CL and BTER. The latter can probably be explained by the generation of isolated nodes instead of degree-1 nodes which

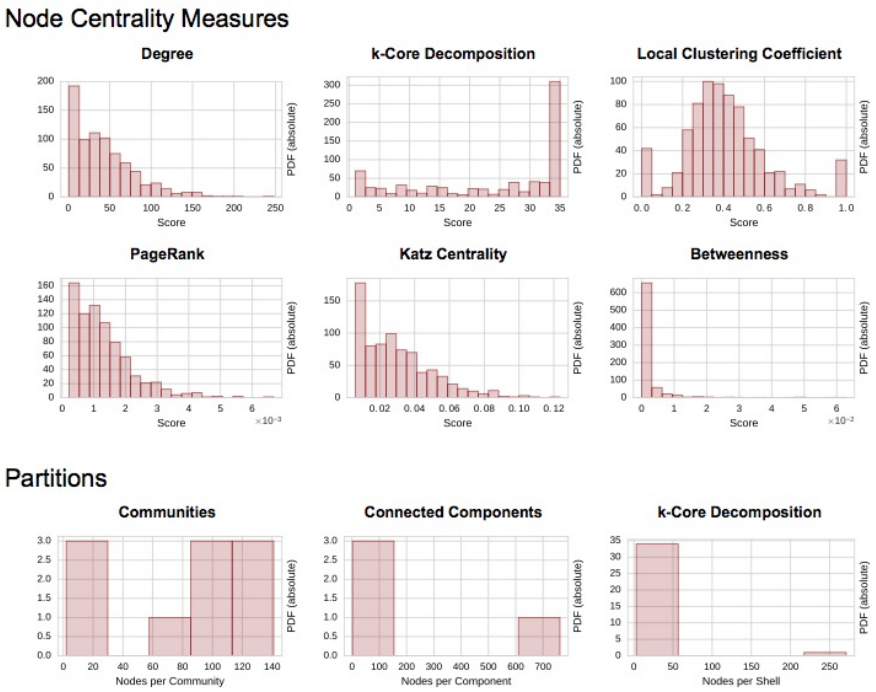

(a) original: `fb-Caltech36`

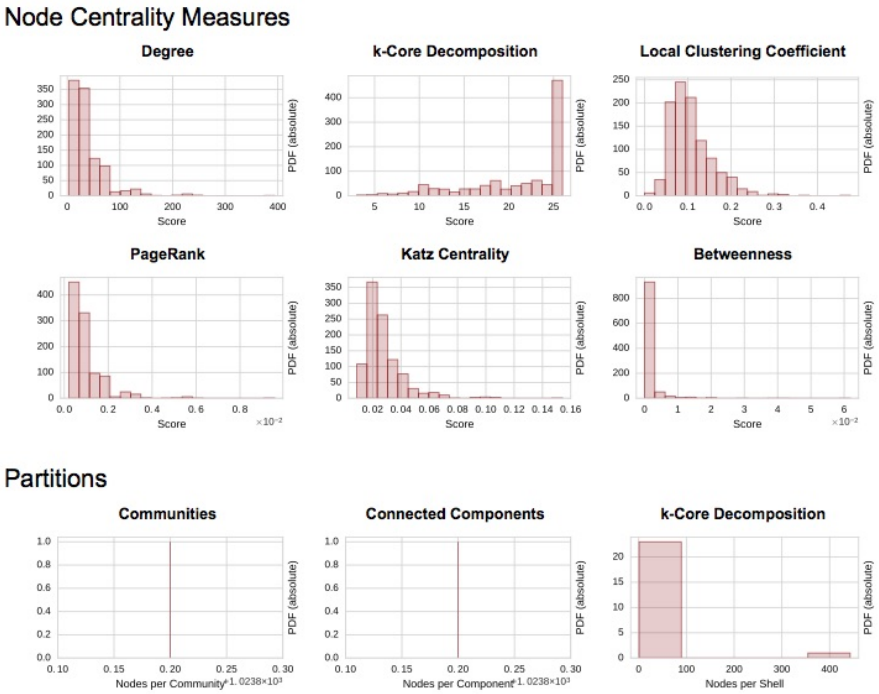

(b) RMAT replica with random initiator matrix

Figure 9: Structure profiles for the evaluation of the kronfit scheme

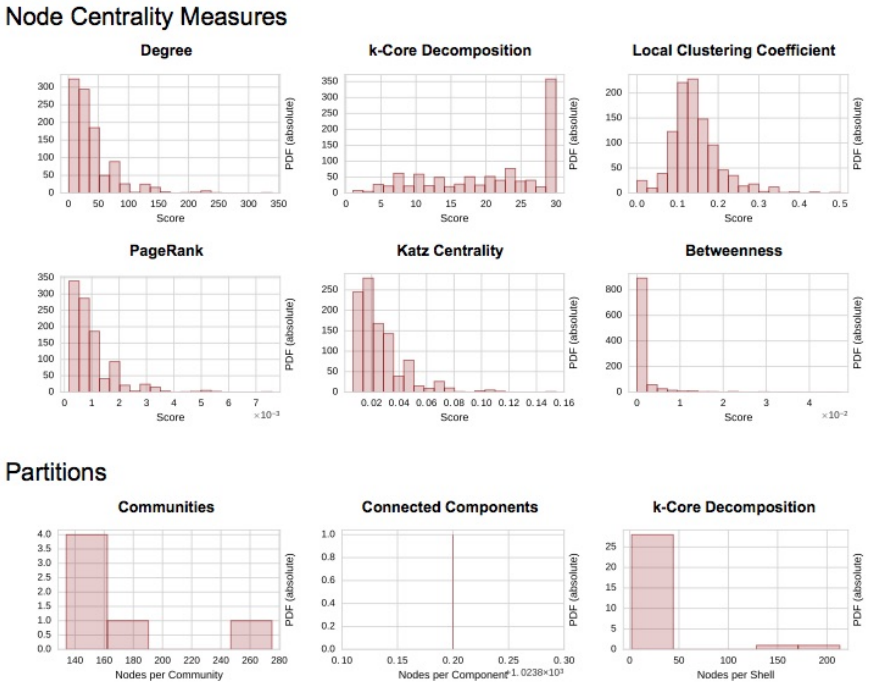

(a) RMAT replica with kronfit-generated initiator matrix

Figure 10: (cont.) Structure profiles for the evaluation of the kronfit scheme

we have already seen in the visualizations before. There is no extreme deviation in terms of the effective diameter, but note that for small world networks, relatively small differences in the effective diameter may indicate significant structural differences. The Barabasi-Albert and Hyperbolic Unit Disk models deviate the most and produce lower diameters. ReCoN, which receives a partition of the original into communities as input, replicates it closely, but all yield different numbers of communities. The Chung-Lu and BTER generators increase the number of communities by a factor of 10 on average which can be explained by the larger number of connected components. Overall, ReCoN emerges with the most accurate replicas from this experiment.

The distributions of node centralities (Fig. 11(b)) compare as follows: All models except ER are capable of producing skewed degree distributions, with CL, ESMC, BTER and ReCoN matching the original closely. Closeness is approximately matched by most models, but BA, RMAT and HUDG deviate significantly. The original networks feature a wide range of clustering coefficients, and only the BTER model, which receives explicitly this distribution, matches them exactly. For HUDG, clustering is extremely artificially high with a median close to 0.9, while LFR produces an unrealistically large variance. For the k-core numbers, random graphs are clearly outliers. RMAT, CL, ESMC, BTER and ReCoN match well, while the very narrow distributions of the others point to a lack of differentiated k-core structure. Only small variations exist with respect to PageRank, and HUDG and LFR have strong deviations in terms of betweenness centrality, but interpretation is not straightforward in this case. In summary, extreme deviations in centrality score distributions clearly give away the artificiality of some synthetic graphs, such as the clustering coefficient for HUDG or the degree distribution for ER. Other differences are more subtle, but possibly relevant. BTER replicates centralities most accurately.

The results on a different set of networks (Tab. 2) are presented in Fig. 12. A notable difference is the extreme variance of clustering coefficients in the original set, which ReCoN cannot replicate. Again ReCoN performs well for the majority of properties.

Fig. 13 shows results of a repetition of the experiment with a scaling factor of 4. All models except RMAT achieve the targeted edge factor of $m' = 4 \cdot m$. The number of components is unrealistically increased by RMAT and CL. For the effective diameter, small relative differences

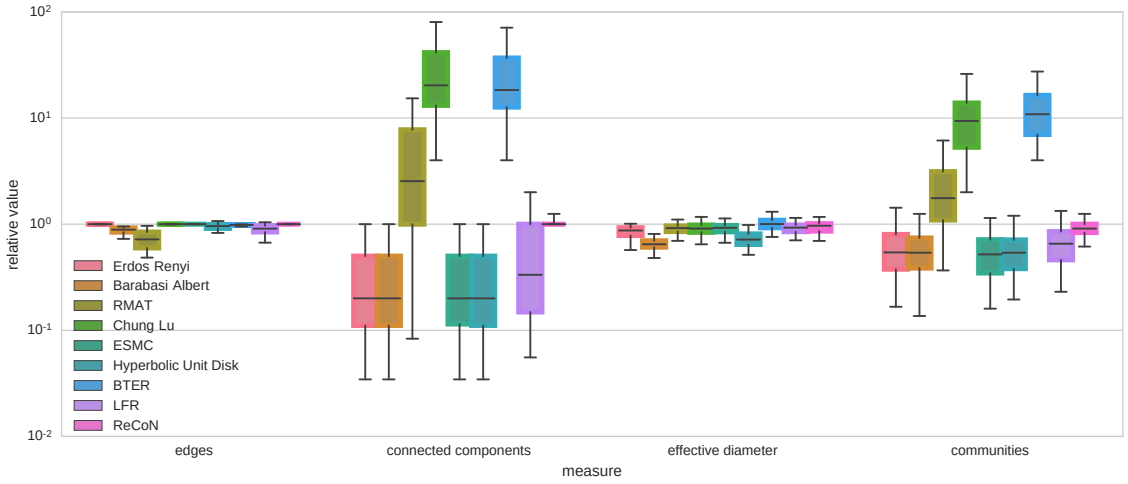

(a) relative difference of scalar network properties

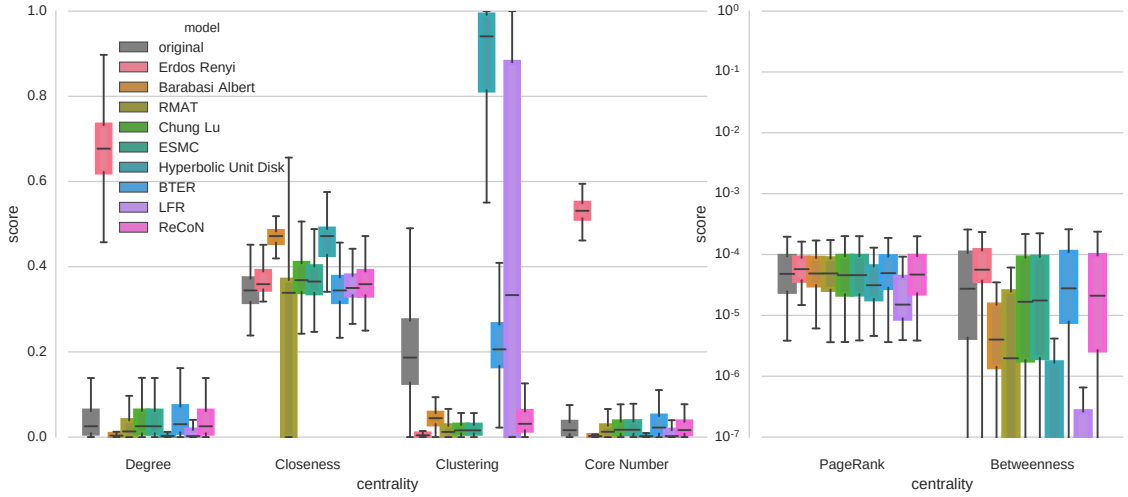

(b) distribution of centrality scores

Figure 11: Structure replication of Facebook networks

Network	**Type**	n	m
email-Enron	email communication	36691	183830
PGPgiantcompo	PGP web of trust	10680	24316
as-22july06	internet topology	22963	48436
hep-th	scientific coauthorship	8361	15751
dolphins	animal social network	62	159
power	power grid	4941	6594
cnr-2000	web graph	325557	2738969

Table 2: Additional networks used. The network email-Enron has been taken from the Stanford Large Network Dataset Collection [42], all other networks are from the clustering instances of the 10th DIMACS implementation challenge [4].

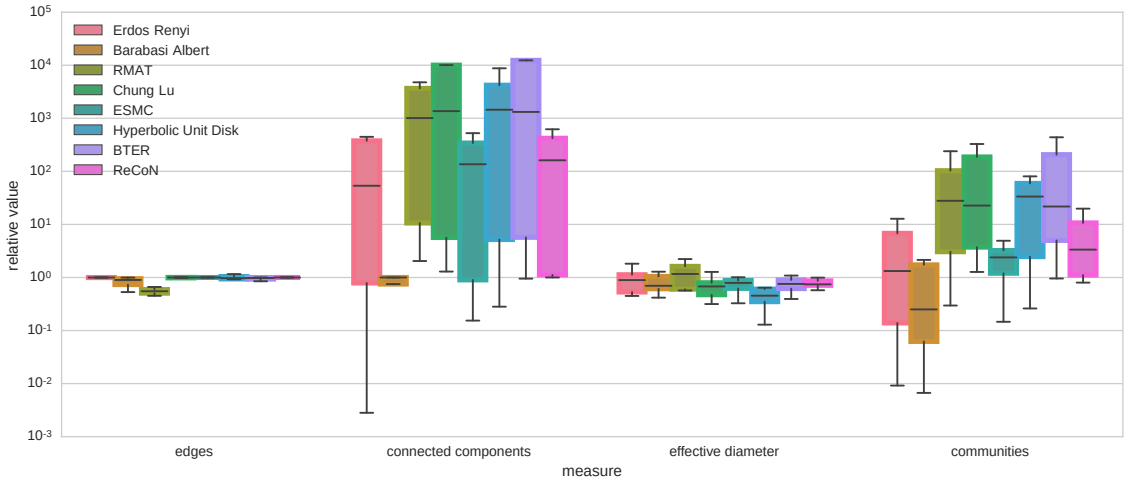

(a) relative difference of scalar network properties

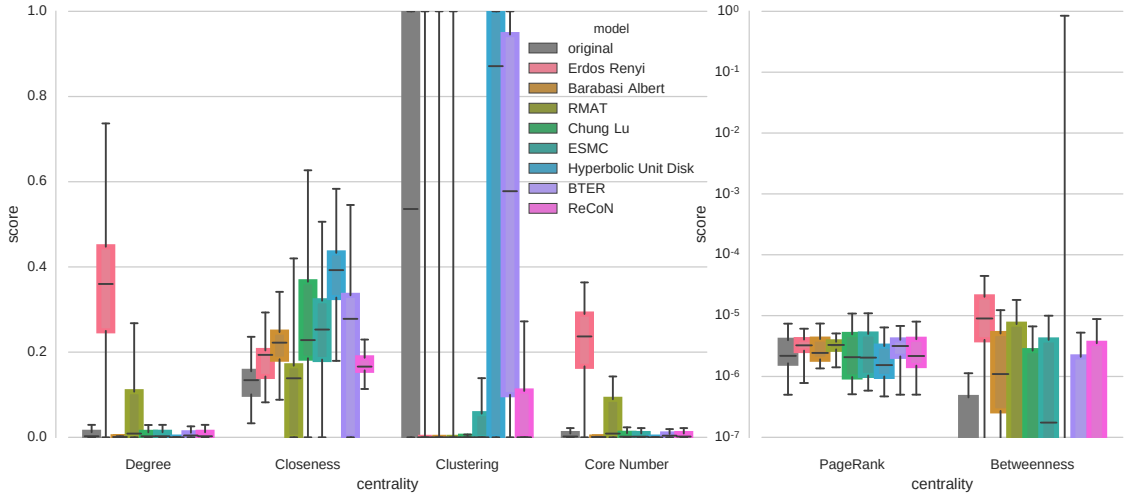

(b) distribution of centrality scores

Figure 12: Structure replication of diverse set of networks (Tab. 2)

matter: BA and HUDG model tend to create smaller worlds than reality. ReCoN produces a remarkably exact match, considering that the generator does not explicitly target the diameter. It does however target a higher number of communities, which is desired and achieved. LFR keeps the number of communities constant on average, while many other models produce fewer communities than the originals. The relative differences in the distributions of centralities are qualitatively equivalent to those in Fig. 11(b).

References

[1] William Aiello, Fan Chung, and Linyuan Lu. A random graph model for massive graphs. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 171–180. Acm, 2000.

[2] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002.

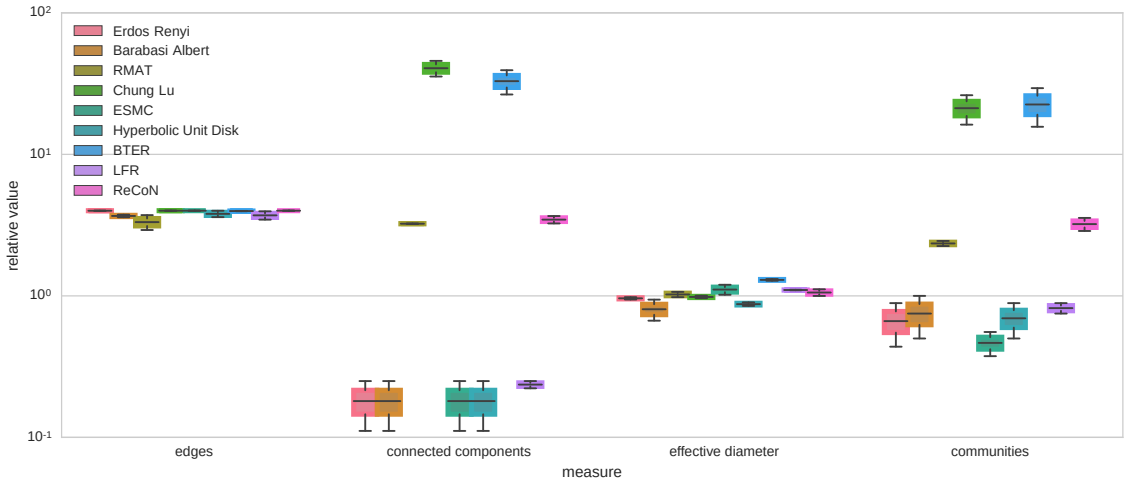

(a) relative difference of scalar network properties

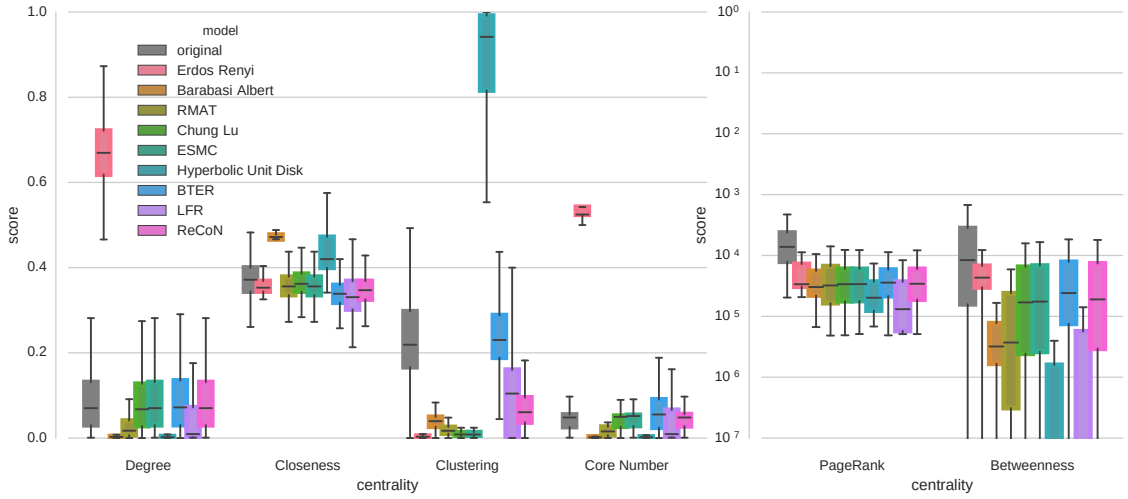

(b) distribution of centrality scores

Figure 13: Structure replication of Facebook networks with scaling factor 4

[3] R. Albert and A.L. Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

[4] David A Bader, Henning Meyerhenke, Peter Sanders, Christian Schulz, Andrea Kappes, and Dorothea Wagner. Benchmarking for graph clustering and partitioning. In *Encyclopedia of Social Network Analysis and Mining*, pages 73–82. Springer, 2014.

[5] S. Bansal, S. Khandelwal, and L.A. Meyers. Exploring biological network structure with clustered random networks. *BMC Bioinformatics*, 10(1):405, 2009.

[6] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, October 1999.

[7] Vladimir Batagelj and Ulrik Brandes. Efficient generation of large random networks. *Physical Review E*, 71(3):036113, 2005.

[8] P. M. Binder. Frustration in complexity. *Science*, 322:323, 2008.

[9] Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. Complex networks: Structure and dynamics. *Physics reports*, 424(4):175–308, 2006.

[10] A. Brandt. Multiscale scientific computation: Review 2001. In T. Barth, R. Haimes, and T. Chan, editors, *Multiscale and Multiresolution Methods (Proceedings of the Yosemite Educational Symposium, October 2000)*. Springer-Verlag, 2001.

[11] A. Brandt and D. Ron. Chapter 1 : Multigrid solvers and multilevel optimization strategies. In J. Cong and J. R. Shinnerl, editors, *Multilevel Optimization and VLSICAD*. Kluwer, 2003.

[12] J. M. Brase and D. L. Brown. Modeling, simulation and analysis of complex networked systems - a program plan. *Lawrence Livermore National Laboratory*, 2009.

[13] W. L. Briggs, V. Henson, and S. F. McCormick. *A multigrid tutorial: second edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2000.

[14] Aydın Buluç, Henning Meyerhenke, Ilya Safro, Peter Sanders, and Christian Schulz. Recent advances in graph partitioning. In *Algorithm Engineering*, pages 117–158. Springer, 2016.

[15] Guido Caldarelli and Alessandro Vespignani. *Large scale structure and dynamics of complex networks*. World Scientific, 2007.

[16] J.M. Carlson and J. Doyle. Complexity and robustness. *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl 1):2538, 2002.

[17] Deepayan Chakrabarti and Christos Faloutsos. Graph mining: Laws, generators, and algorithms. *ACM Computing Surveys (CSUR)*, 38(1):2, 2006.

[18] Deepayan Chakrabarti, Yiping Zhan, and Christos Faloutsos. R-MAT: A recursive model for graph mining. *Computer Science Department*, page 541, 2004.

[19] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[20] Naga Shailaja Dasari, Desh Ranjan, and Mohammad Zubair. ParK: An efficient algorithm for k-core decomposition on multicore processors. In *2014 IEEE International Conference on Big Data, Big Data 2014, Washington, DC, USA, October 27-30, 2014*, pages 9–16. IEEE, 2014.

[21] D.M. Dunlavy, B. Hendrickson, and T.G. Kolda. Mathematical challenges in cybersecurity. Technical report, Sandia National Laboratory, 2009.

[22] P. Erdos and A. Renyi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5:17–61, 1960.

[23] S. Eubank, H. Guclu, V.S.A. Kumar, M.V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184, 2004.

[24] Santo Fortunato. Benchmark graphs to test community detection algorithms. `https://sites.google.com/site/santofortunato/inthepress2`.

[25] Robert Geisberger, Peter Sanders, and Dominik Schultes. Better approximation of betweenness centrality. In *ALENEX*, pages 90–100. SIAM, 2008.

[26] Roland Glantz, Henning Meyerhenke, and Christian Schulz. Tree-based coarsening and partitioning of complex networks. *Journal of Experimental Algorithmics (JEA)*, 21(1):1–6, 2016.

[27] Anna Goldenberg, Alice X Zheng, Stephen E Fienberg, and Edoardo M Airoldi. A survey of statistical network models. *Foundations and Trends® in Machine Learning*, 2(2):129–233, 2010.

[28] A. Gutfraind, L.A. Meyers, and I. Safro. MUSKETEER: Multiscale entropic network generator. `https://people.cs.clemson.edu/~isafro/musketeer/index.html`, 2012.

[29] Alexander Gutfraind, Ilya Safro, and Lauren Ancel Meyers. Multiscale network generation. In *Information Fusion (Fusion), 2015 18th International Conference on*, pages 158–165. IEEE, 2015.

[30] Michael Hamann, Gerd Lindner, Henning Meyerhenke, Christian L. Staudt, and Dorothea Wagner. Structure-preserving sparsification methods for social networks. *Social Netw. Analys. Mining*, 6(1):22:1–22:22, 2016.

[31] S. Itzkovitz, R. Levitt, N. Kashtan, R. Milo, M. Itzkovitz, and U. Alon. Coarse-graining and self-dissimilarity of complex networks. *Physical Review E*, 71(1):016127, 2005.

[32] Matthew James Keeling and Pejman Rohani. *Modeling infectious diseases in humans and animals*, volume 47. Princeton University Press, 2008.

[33] Tamara G Kolda, Ali Pinar, Todd Plantenga, and C Seshadhri. A scalable generative graph model with community structure. *arXiv preprint arXiv:1302.6636*, 2013.

[34] P. L. Krapivsky and S. Redner. Organization of growing random networks. *Phys. Rev. E*, 63(066123), 2001.

[35] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Physical Review E*, 82:036106, Sep 2010.

[36] Jérôme Kunegis. Konect: the koblenz network collection. In *Proceedings of the 22nd international conference on World Wide Web companion*, pages 1343–1350. International World Wide Web Conferences Steering Committee, 2013.

[37] Andrea Lancichinetti and Santo Fortunato. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Physical Review E*, 80(1):016118, 2009.

[38] Andrea Lancichinetti, Santo Fortunato, and Filippo Radicchi. Benchmark graphs for testing community detection algorithms. *Physical Review E*, 78(4):046110, 2008.

[39] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *The Journal of Machine Learning Research*, 11:985–1042, 2010.

[40] J. Leskovec, K.J. Lang, A. Dasgupta, and M.W. Mahoney. Statistical properties of community structure in large social and information networks. In *Proceeding of the 17th international conference on World Wide Web*, pages 695–704. ACM, 2008.

[41] Jure Leskovec and Christos Faloutsos. Scalable modeling of real graphs using kronecker multiplication. In *Proc. 24th Intl. Conference on Machine learning*, pages 497–504. ACM, 2007.

[42] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection. `http://snap.stanford.edu/data`, June 2014.

[43] Jure Leskovec and Rok Sosič. SNAP: A general purpose network analysis and graph mining library in C++. `http://snap.stanford.edu/snap`, June 2014.

[44] Sven Leyffer and Ilya Safro. Fast response to infection spread and cyber attacks on large-scale networks. *Journal of Complex Networks*, 1(2):183–199, 2013.

[45] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. *ACM SIGCOMM Computer Communication Review*, 36(4):135–146, 2006.

[46] Alberto Medina, Ibrahim Matta, and John Byers. On the origin of power laws in internet topologies. Technical report, Boston University, Boston, MA, USA, 2000.

[47] Henning Meyerhenke, Martin Nöllenburg, and Christian Schulz. Drawing large graphs by multilevel maxent-stress optimization. In Emilio Di Giacomo and Anna Lubiw, editors, *Graph Drawing and Network Visualization - 23rd International Symposium, GD 2015, Los Angeles, CA, USA, September 24-26, 2015, Revised Selected Papers*, volume 9411 of *Lecture Notes in Computer Science*, pages 30–43. Springer, 2015.

[48] Henning Meyerhenke, Peter Sanders, and Christian Schulz. Partitioning (hierarchically clustered) complex networks via size-constrained graph clustering. *J. Heuristics*, 22(5):759–782, 2016.

[49] L.A. Meyers, B. Pourbohloul, M.E.J. Newman, D.M. Skowronski, and R.C. Brunham. Network theory and sars: predicting outbreak diversity. *Journal of theoretical biology*, 232(1):71–81, 2005.

[50] C.G.M. Mihail and E. Zegura. The markov chain simulation method for generating connected power law random graphs. In *Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments*, volume 111, page 16. Society for Industrial Mathematics, 2003.

[51] R. Milo, N. Kashtan, S. Itzkovitz, M. E. J. Newman, and U. Alon. On the uniform generation of random graphs with prescribed degree sequences. *eprint arXiv:cond-mat/0312028*, December 2003.

[52] Enys Mones, Lilla Vicsek, and Tamas Vicsek. Hierarchy measure for complex networks. *PLoS ONE*, 7(3):e33799, 03 2012.

[53] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

[54] Mark Newman. *Networks: an introduction*. Oxford University Press, 2010.

[55] G. Palla, L. Lovász, and T. Vicsek. Multifractal network generator. *Proceedings of the National Academy of Sciences*, 107(17):7640, 2010.

[56] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.

[57] Christopher R Palmer, Phillip B Gibbons, and Christos Faloutsos. Anf: A fast and scalable tool for data mining in massive graphs. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 81–90. ACM, 2002.

[58] John J Potterat, L Phillips-Plummer, Stephen Q Muth, RB Rothenberg, DE Woodhouse, TS Maldonado-Long, HP Zimmerman, and JB Muth. Risk network structure in the early epidemic phase of hiv transmission in colorado springs. *Sexually transmitted infections*, 78(suppl 1):i159–i163, 2002.

[59] D. Ron, I. Safro, and A. Brandt. Relaxation-based coarsening and multiscale graph organization. *Multiscale Modeling & Simulation*, 9(1):407–423, 2011.

[60] Wolfgang E. Schlauch, Emőke Ágnes Horvát, and Katharina A. Zweig. Different flavors of randomness: comparing random graph models with fixed degree sequences. *Social Network Analysis and Mining*, 5(1):1–14, 2015.

[61] C. Seshadhri, Tamara G. Kolda, and Ali Pinar. Community structure and scale-free collections of Erdös-Renyi graphs. December 2011.

[62] Tom AB Snijders. The statistical evaluation of social network dynamics. *Sociological methodology*, 31(1):361–395, 2001.

[63] Christian L. Staudt. *Algorithms and Software for the Analysis of Large Complex Networks*. PhD thesis, Karlsruhe Institute of Technology, June 2016.

[64] Christian L Staudt, Michael Hamann, Ilya Safro, Alexander Gutfraind, and Henning Meyerhenke. Generating scaled replicas of real-world complex networks. In *International Workshop on Complex Networks and their Applications*, pages 17–28. Springer, 2016.

[65] Christian L. Staudt and Henning Meyerhenke. Engineering parallel algorithms for community detection in massive networks. *IEEE Trans. on Parallel and Distributed Systems*, 27(1):171–184, 2016.

[66] Christian L. Staudt, Aleksejs Sazonovs, and Henning Meyerhenke. NetworKit: A tool suite for large-scale complex network analysis. *Network Science*, 4(4):508530, 2016.

[67] Sandia National Laboratories Tamara G. Kolda, Ali Pinar. Feastpack. `http://www.sandia.gov/~tgkolda/feastpack/`.

[68] Amanda L Traud, Peter J Mucha, and Mason A Porter. Social structure of facebook networks. *Physica A: Statistical Mechanics and its Applications*, 391(16):4165–4180, 2012.

[69] U. Trottenberg and A. Schuller. *Multigrid*. Academic Press, Inc., 2001.

[70] Fabien Viger and Matthieu Latapy. Random generation of large connected simple graphs with prescribed degree distribution. In *11th International Conference on Computing and Combinatorics. Kunming, Yunnan, Chine*, 2005.

[71] Moritz von Looz, Henning Meyerhenke, and Roman Prutkin. Generating random hyperbolic graphs in subquadratic time. In *Algorithms and Computation - 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9-11, 2015, Proceedings*, pages 467–478, 2015.

[72] D J Watts and S H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–2, June 1998.

[73] D.H. Wolpert and W. Macready. Using self-dissimilarity to quantify complexity. *Complexity*, 12(3):77–85, 2007.